NeuralTrust

The Platform For AI and Agent Security

NeuralTrust uncovers vulnerabilities, blocks attacks, monitors performance, and ensures regulatory compliance — everything enterprises need to scale AI and Agents with confidence

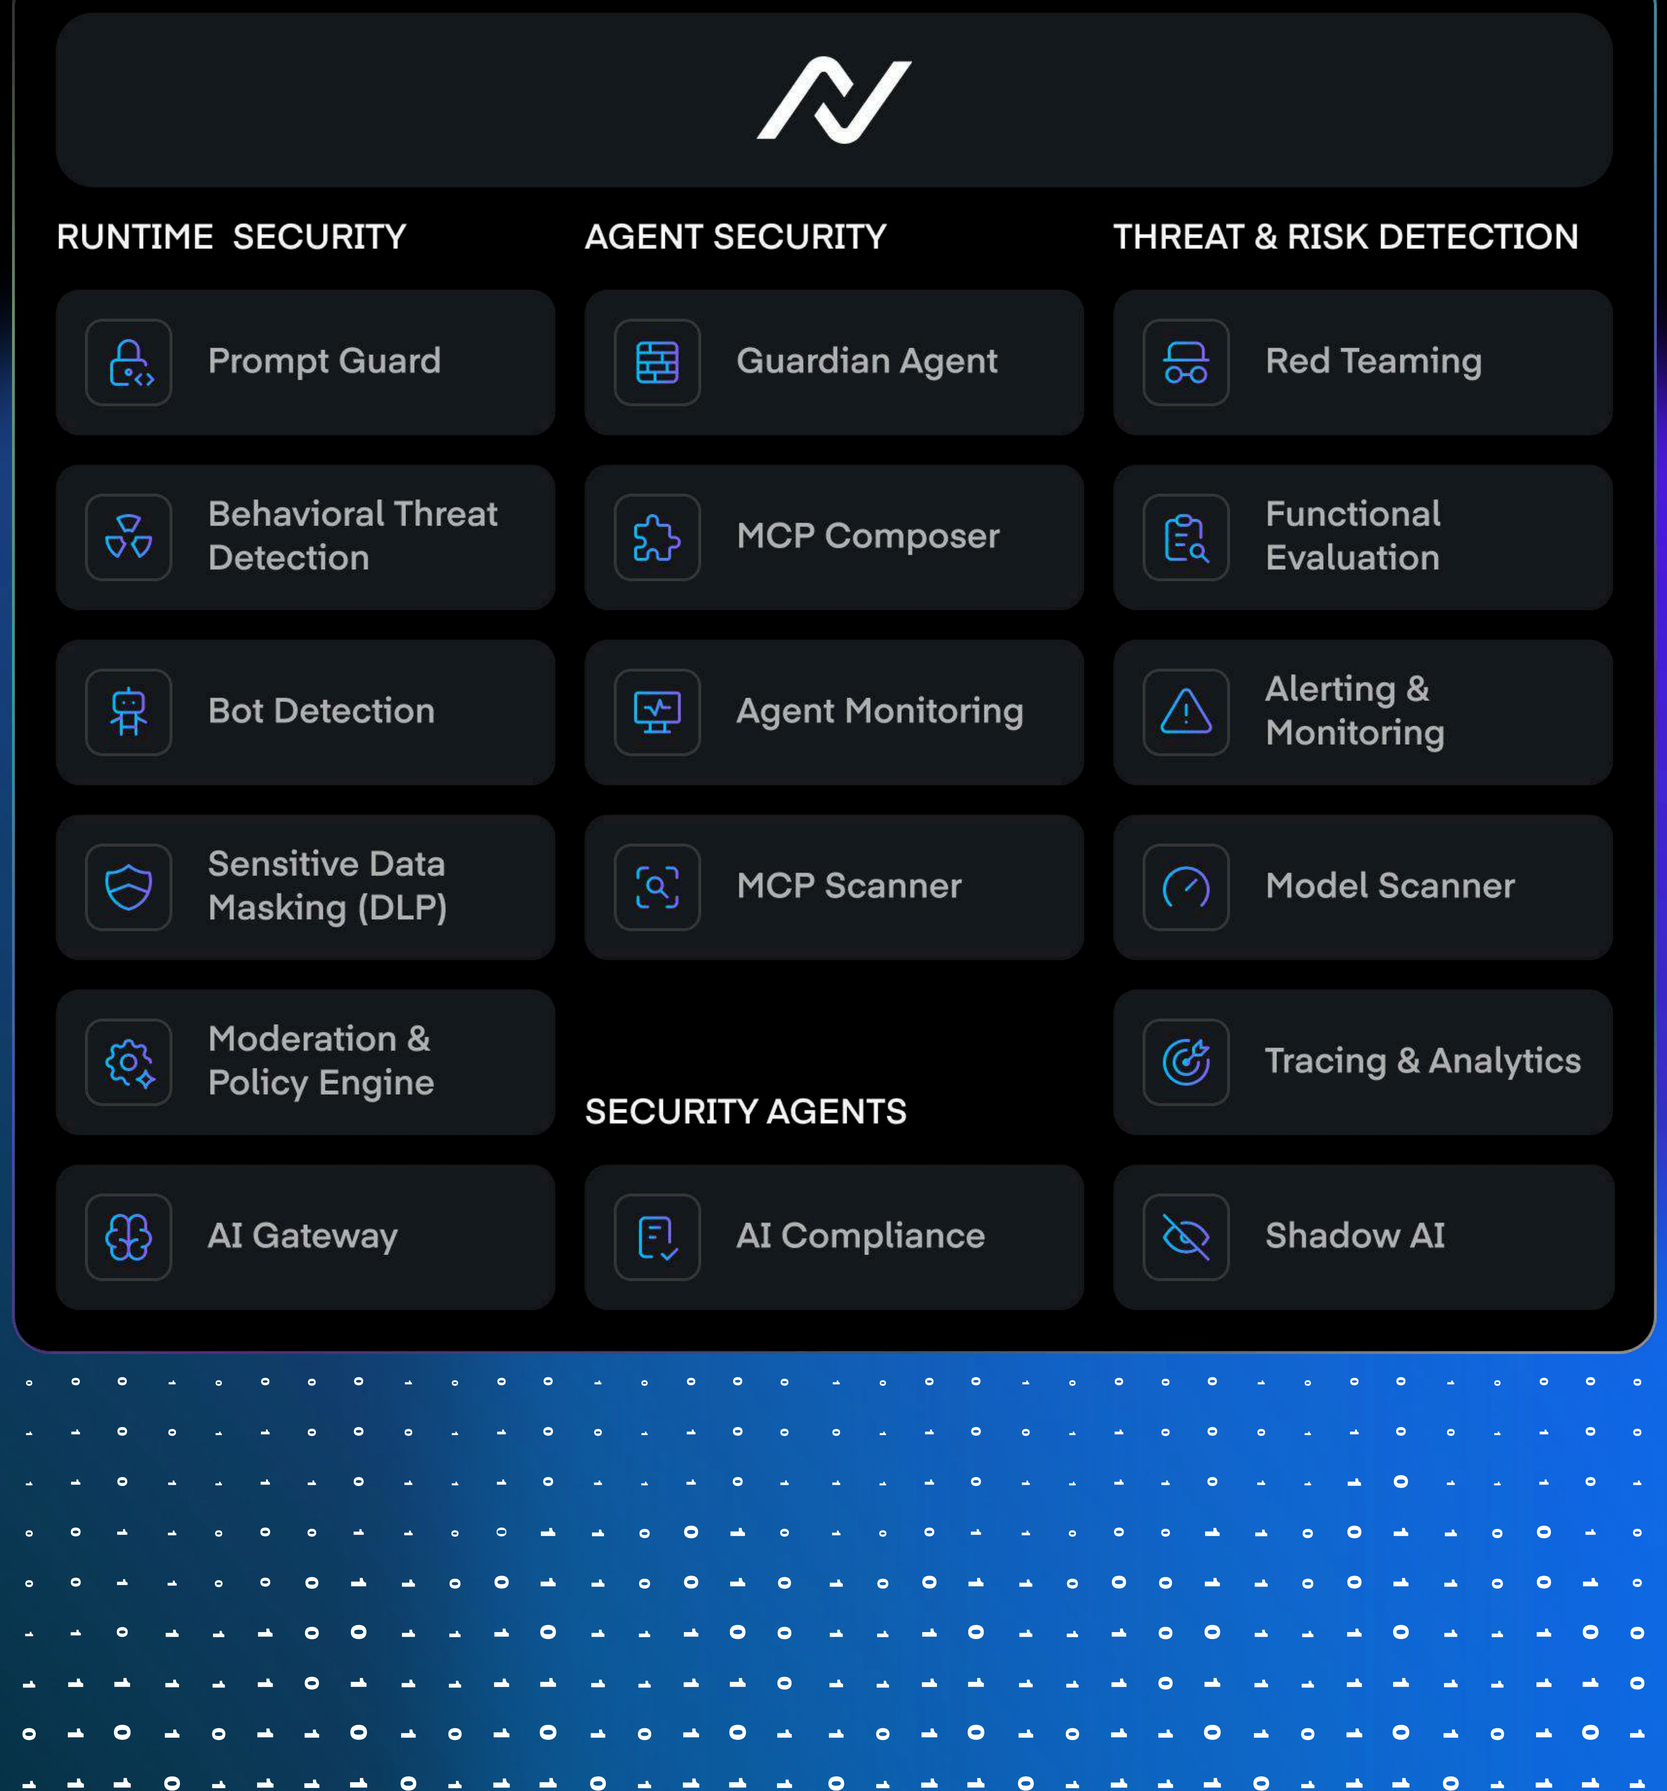

NeuralTrust

Scale Generative AI With Confidence

Generative AI is transforming everything from customer service to internal productivity, but it also introduces new risks: jailbreaks, data leaks, unmonitored AI behavior... As use cases grow, so does the need for security, oversight, and standardized control.

The impact of unsafe AI in numbers

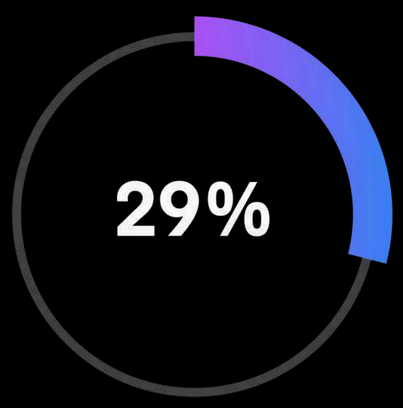

of companies deploying AI have suffered a security breach

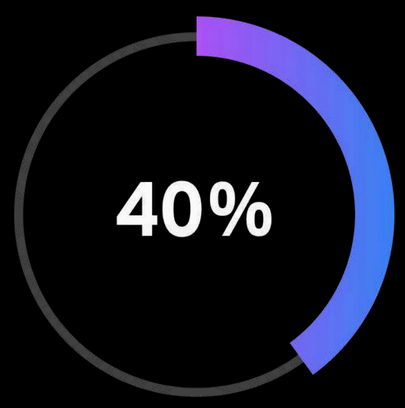

of AI breaches stem from misuse of generative AI

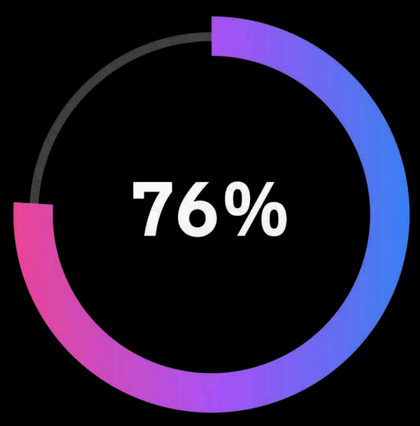

of AI initiatives are not currently secured against cyber risks

Our Key Solutions:

NeuralTrust offers a comprehensive suite of tools to seamlessly implement the mitigation strategies outlined. This section details how our platform can help you tackle each identified vulnerability step-by-step, ensuring your AI systems are secure, reliable, and compliant.

NeuralTrust is the most performant and scalable ecosystem for LLM security and control: comprehensive, powerful, and built for the future.

1. AI Runtime Security

A unified runtime layer that intercepts and sanitizes every LLM request to block prompt injections and unsafe inputs.

2. AI Agent Security

From prompt input to tool execution, NeuralTrust gives you complete control over how autonomous agents behave, and what they can access.

3. Threat & Risk Detection

Continuously test and monitor your AI with adaptive red teaming, real-time alerts, automated vulnerability scans, detailed tracing and analytics.

4. Security Agents

Our AI Agent does the tedious compliance work for you: from regulatory monitoring to policy updates.

TRUSTED BY THE WORLD LEADING COMPANIES

Sabadell · IBERIA · ABANCA · ISDIN · AirEuropa

Capgemini · LUGAPEL · NTT DATA · devoteam · Bentego

NeuralTrust

1. Runtime Security (GAF)

The Generative Application Firewall (GAF) by NeuralTrust is the next-generation security solution built to protect generative AI applications from evolving threats, without compromising performance or flexibility. Unlike traditional application firewalls or basic prompt-level defenses, GAF operates across the full stack: from infrastructure and API layers to runtime behavior and semantic output.

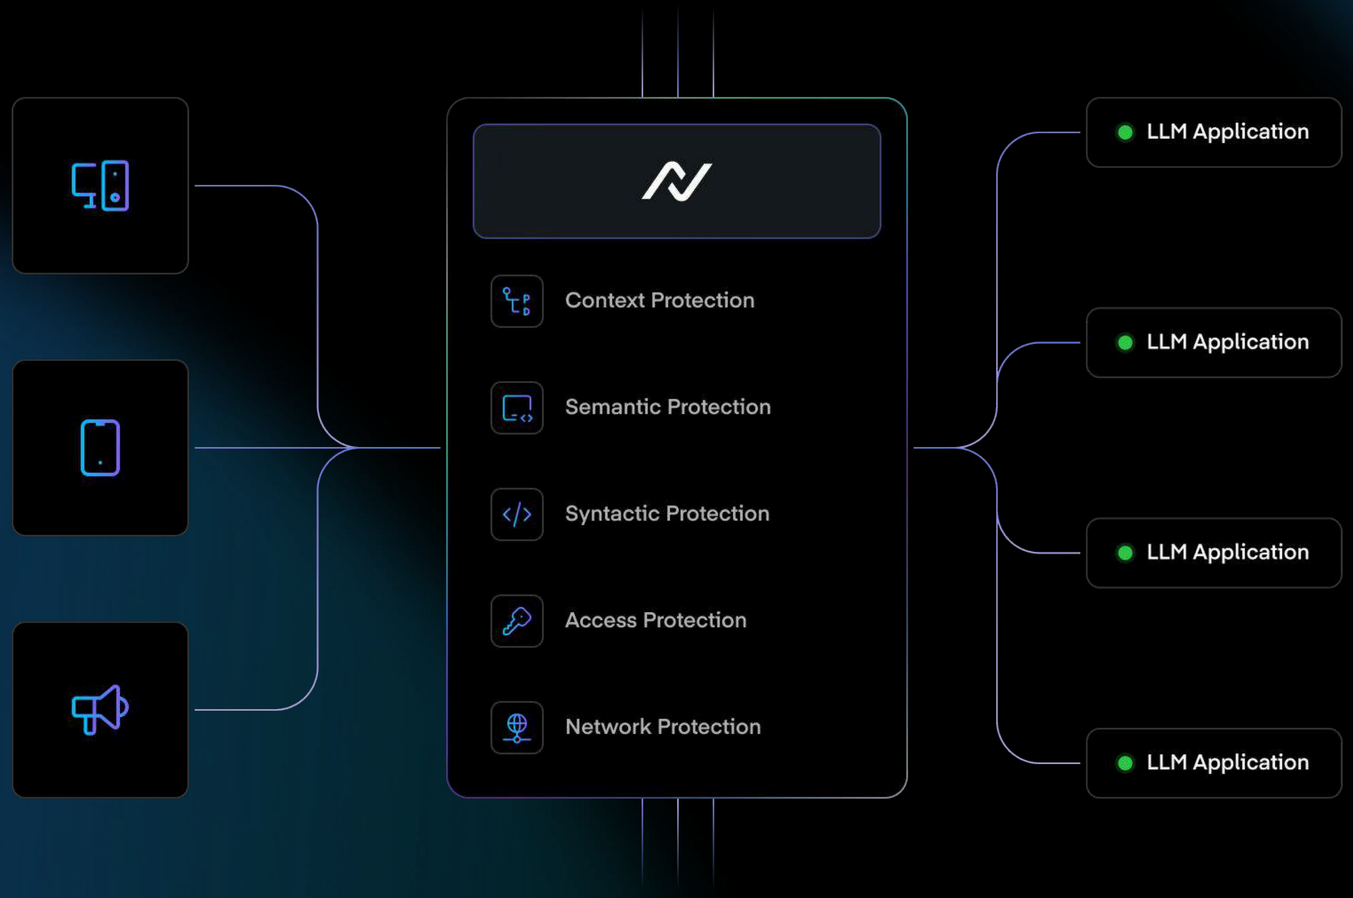

Why a Generative Application Firewall? (GAF)

As generative systems become more powerful and widely adopted, they introduce new risks: prompt injections, unauthorized tool use, multi-turn exploits, and data leakage. GAF mitigates these risks by enforcing dynamic, organization-wide security policies and adopting a zero-trust approach to generative workflows. It delivers real-time behavioral threat detection, automated policy enforcement, and seamless integration with your existing AI infrastructure. This way you can scale AI safely and confidently.

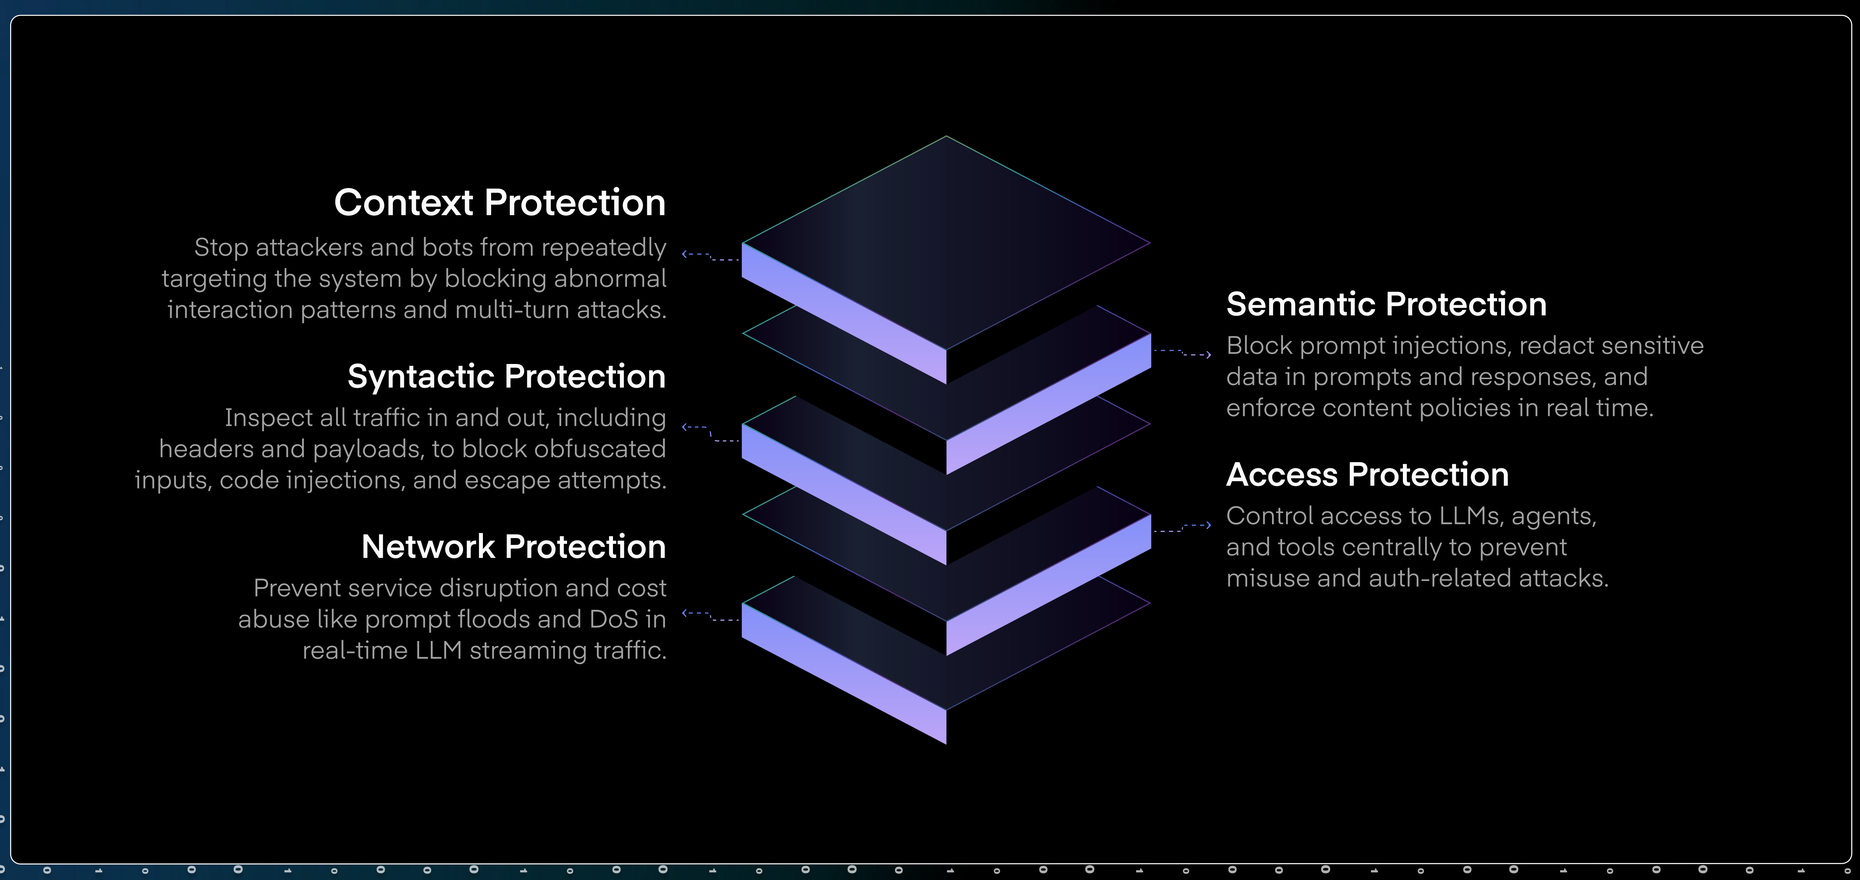

NeuralTrust

Scalability and Efficiency

NeuralTrust delivers a unified, real-time defense layer for your Generative AI deployments, so you can confidently power experiences without fear of misuse, leakage, or abuse.

End-to-end security:

GAF protects every layer of your generative AI stack, from prompt inputs to network traffic, detecting bots, multi-turn attacks, and behavioral threats in real time.

Leading performance:

It achieves industry-best detection accuracy with <10ms latency (GPU) and scales linearly to handle 20,000+ requests per second on commodity hardware.

High ceiling:

Designed for flexibility, GAF supports deep customization, plugin-based extensions, and is available as open source to reduce vendor lock-in.

Platform agnostic:

GAF works across all major LLM providers, supports cloud, on-prem, or hybrid deployments, and integrates easily with SIEMs, auth systems, and enterprise infrastructure.

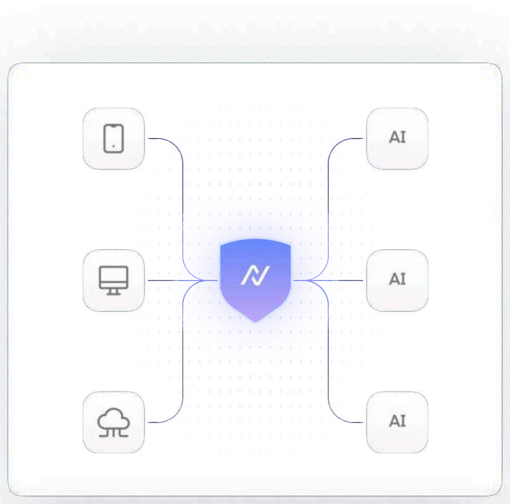

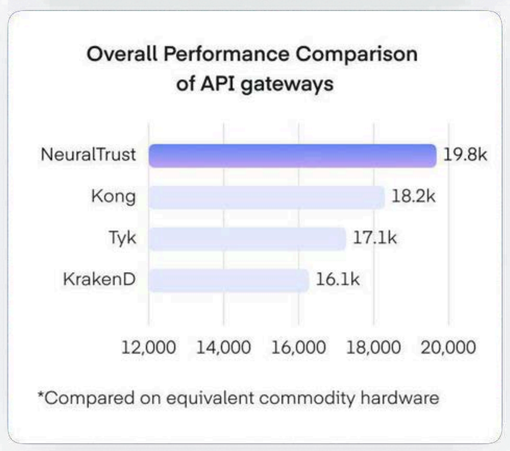

Advanced capabilities

The GAF includes specialized modules to cover every critical risk surface. Whether you're deploying agents, assistants, or full-scale generative platforms, GAF gives you the security foundation to move fast and with confidence.

Prompt Guard	Behavioral Threat Detection	Bot Detection
Monitor usage patterns in real time to identify abnormal, risky, or compromised interactions before they escalate into security incidents.	Monitor usage patterns in real time to identify abnormal, risky, or compromised interactions before they escalate into security incidents.	Detect and block automated traffic, scrapers, and synthetic users before they drain your tokens, hijack your data, or poison your results.
Sensitive Data Masking	**Moderation & Policy Engine**	**AI Gateway**
Automatically detect and redact PII, credentials, and financial information in LLM prompts and responses before it reaches the wrong hands.	Create and enforce moderation policies so you can automatically route flagged content, apply tailored remediations, and oversee reviews in line with your workflows.	Centralize every critical layer of LLM operations into a single control point, enabling unified governance, streamlined integration, and full-stack visibility.

NeuralTrust

2. AI Agent Security

From prompt input to tool execution, NeuralTrust gives you complete control over how autonomous agents behave and what they can access.

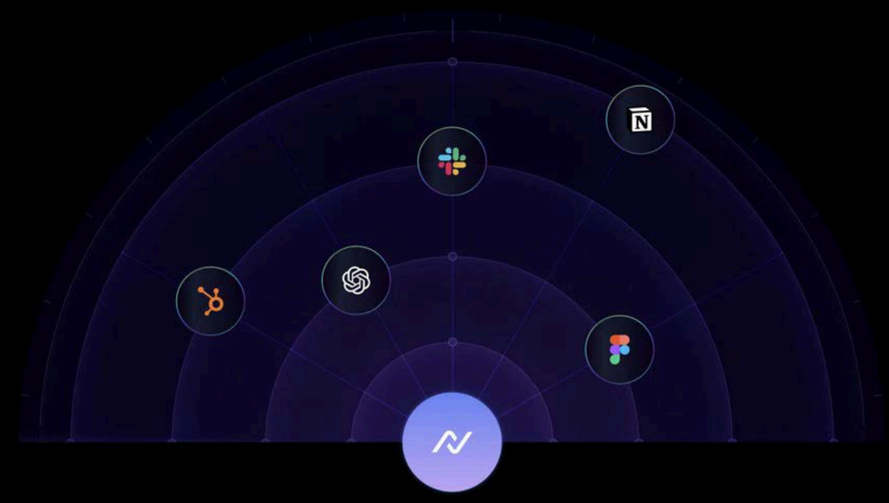

NeuralTrust's Agentic Security Products

1. Guardian Agent:

Secure multi-agent systems and tool-calling workflows against injections, abuse, and unintended actions in real time.

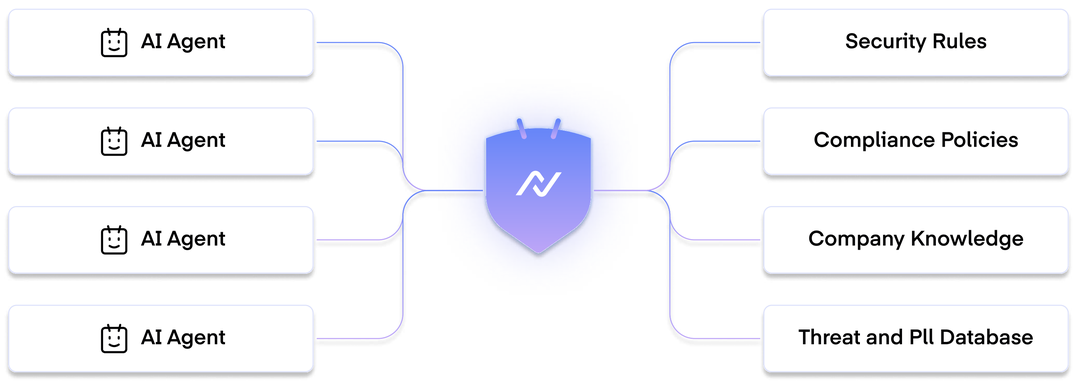

2. Agent Monitoring:

Enforces fine-grained, role-based access controls on your agent control plane (MCP), ensuring tools and data are only invoked by authorized identities.

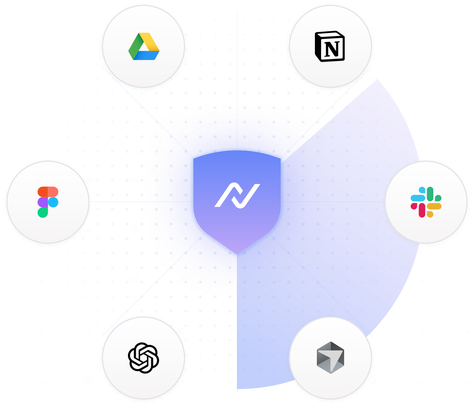

NeuralTrust

3. MCP Composer:

MCP Composer gives your teams complete control over AI Agent interactions with tools and data, with granular permissions for every operation.

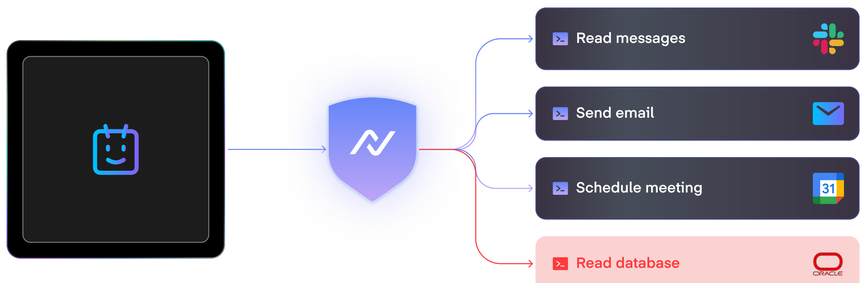

4. MCP Scanner:

Scan and test MCP servers code for threats and vulnerabilities. NeuralTrust ensures MCP servers and tools remain trustworthy as organizations scale agents.

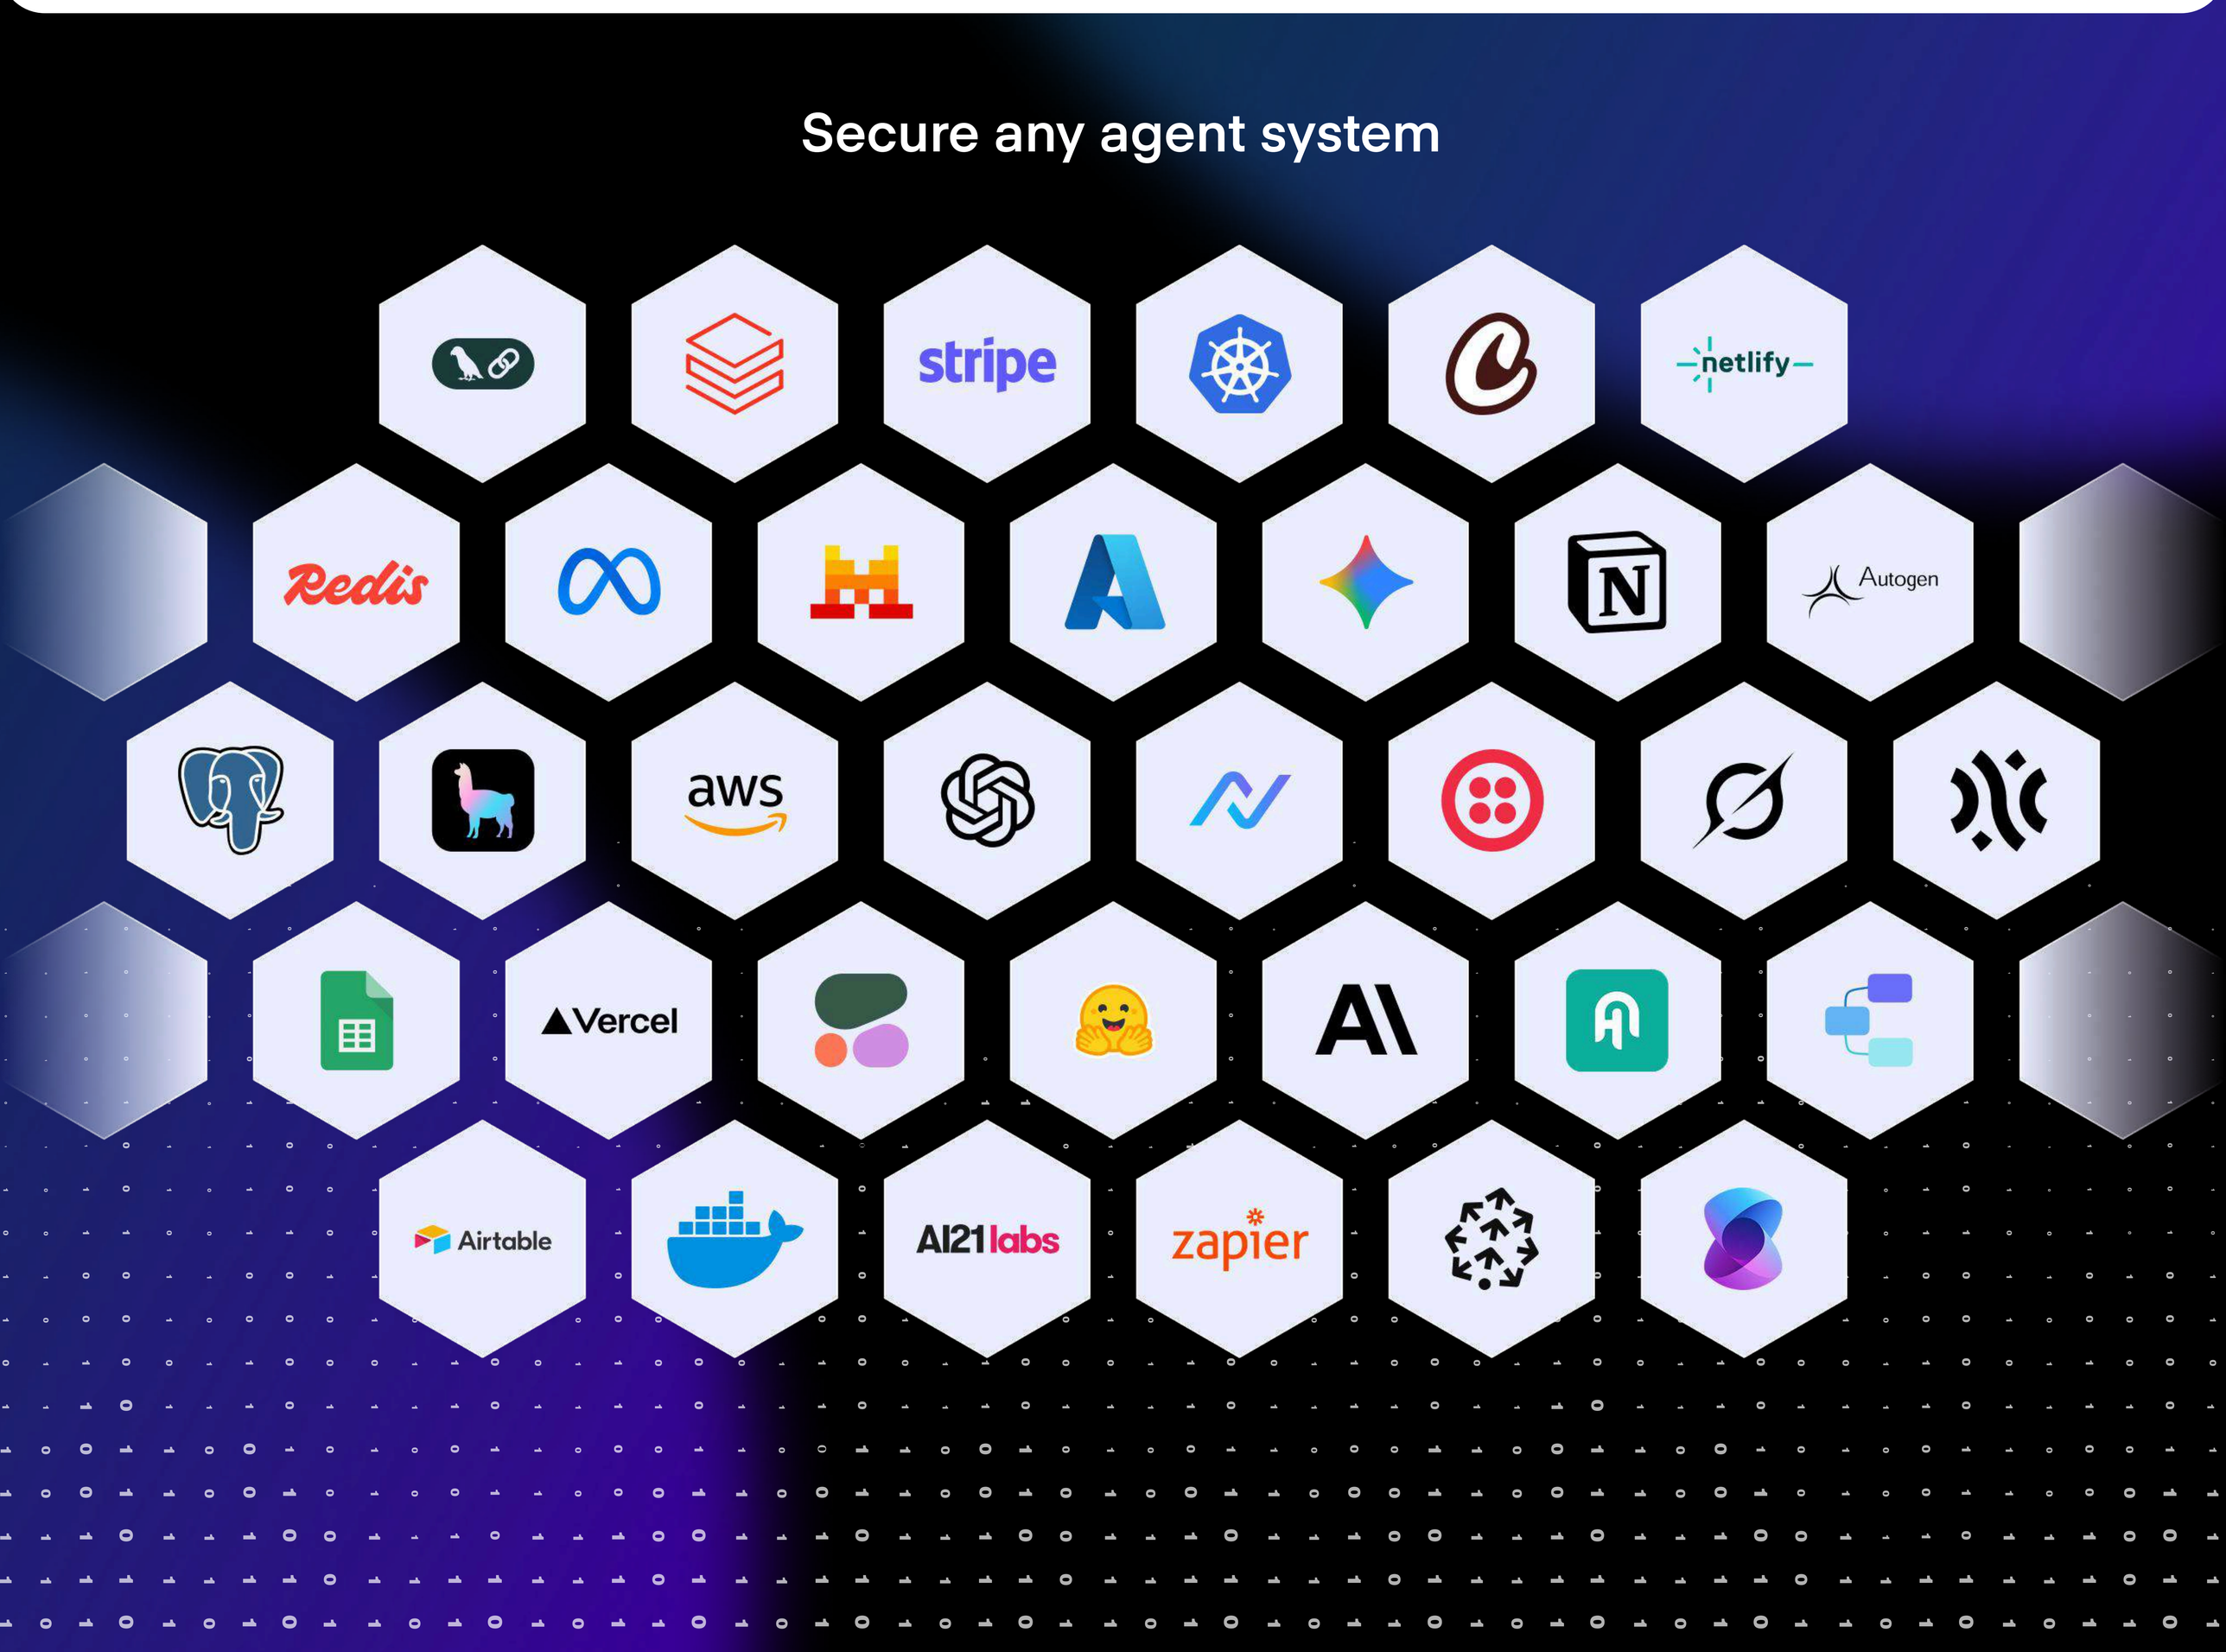

NeuralTrust

3. Threat & Risk Detection

Continuously test and monitor your AI with adaptive red teaming, real-time alerts, automated vulnerability scans, detailed tracing, and conversational analytics. Discover the trusted solution for security and AI teams.

Key Features of NeuralTrust's Observability Platform

1. Automated red teaming for generative AI

Assess your Gen AI apps for vulnerabilities, hallucinations, and errors before they impact your users with a testing platform built for robustness and efficiency.

2. Functional Evaluation

Validate how your GenAI applications behave under different conditions and ensure they meet the standards your users and regulators expect.

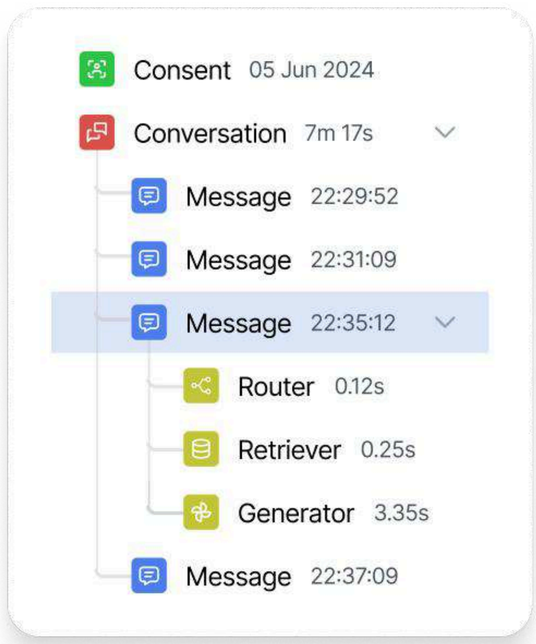

3. Real-time security alerting

Detect threats as they happen with real-time alerts and deep visibility into LLM behavior. Monitor conversations, flag anomalies, and ensure security and compliance at every step.

4. Model Code Scanner

Secure your AI supply chain by identifying malicious code, hidden vulnerabilities, and unsafe agent behavior before deployment.

5. Tracing and Analytics

Identify security flaws, unsafe configurations, and data leaks in your LLM pipeline before deployment.

NeuralTrust

6. Prevent Shadow AI

Track trends, analyze performance, and uncover insights from real LLM conversations with visual dashboards.

4. Security Agents

Our AI Agent does the tedious compliance work for you: from regulatory monitoring to policy updates.

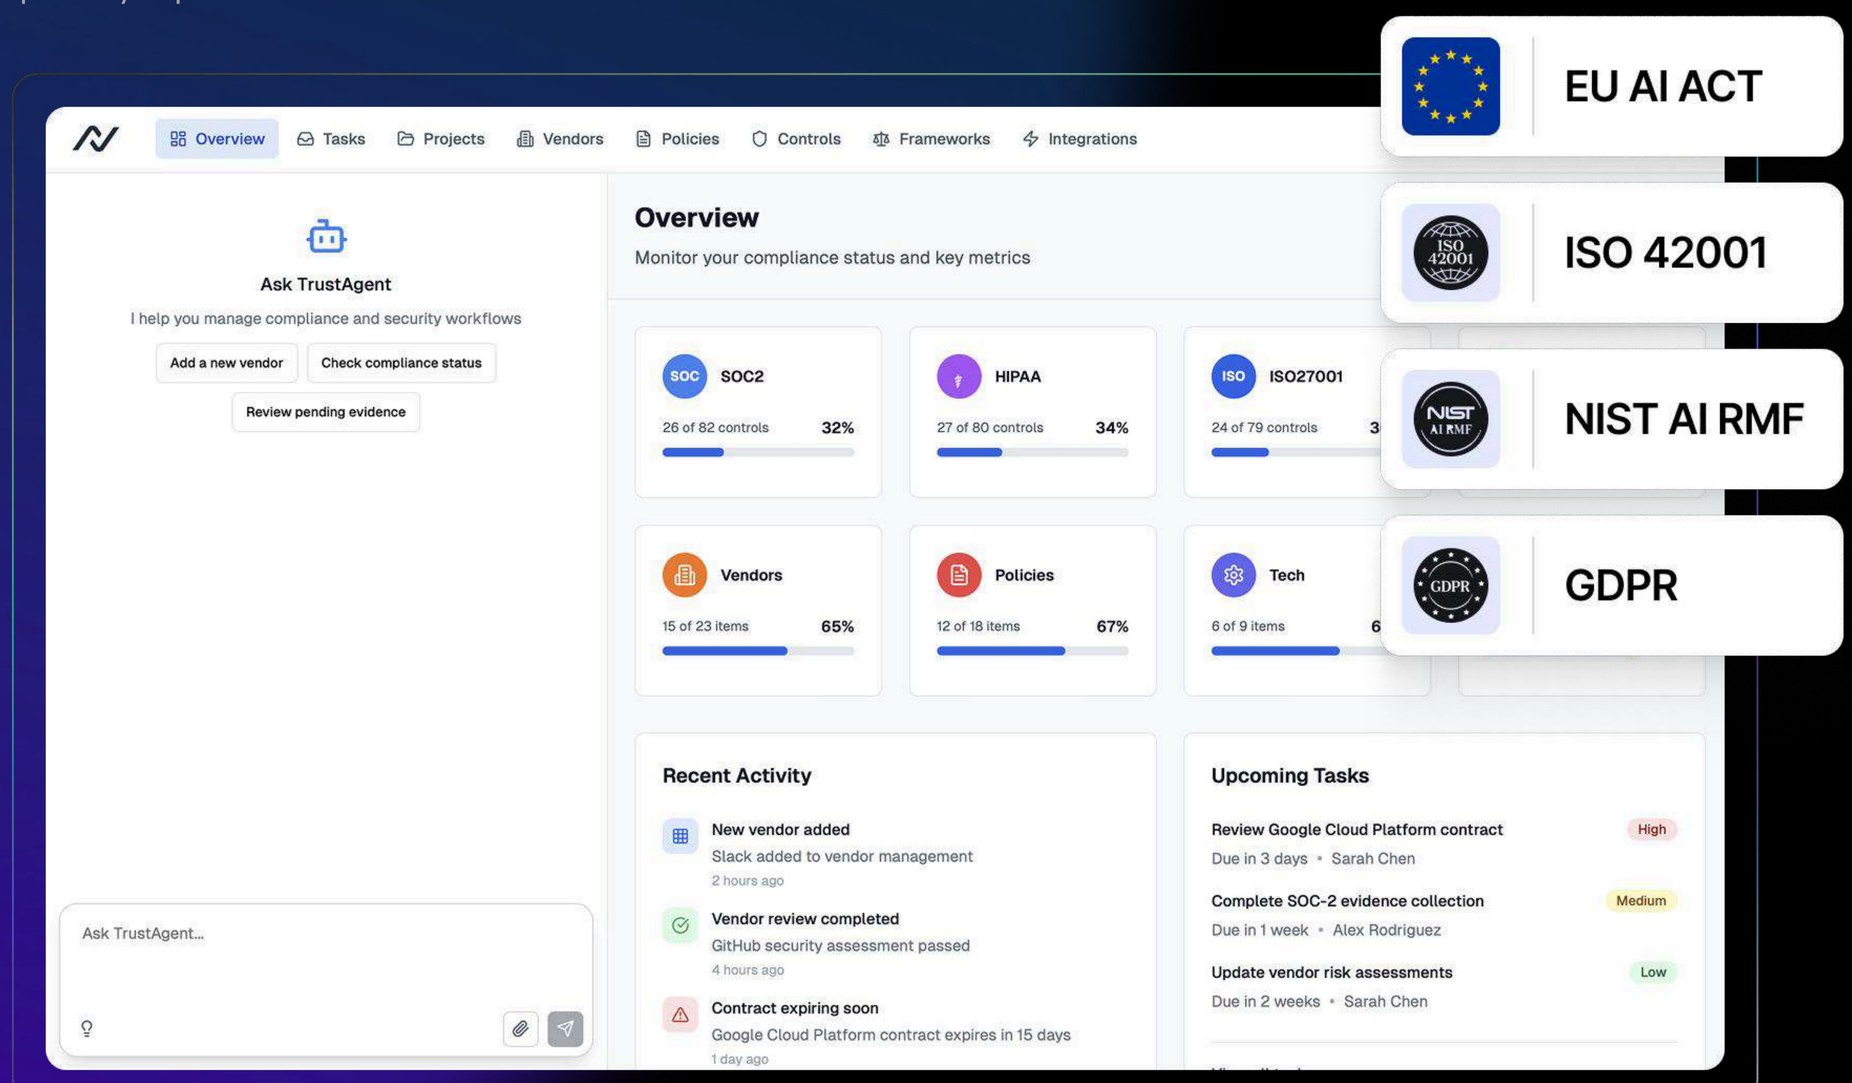

Policy Maker

Create, customize, and manage AI-specific policies based on the EU AI Act, NIST AI RMF, ISO 42001, and more.

Evidence & Controls

Attach technical controls to each policy, and automatically collect audit-ready evidences.

Framework Mapping

Map your policies and controls to compliance frameworks like OWASP, MITRE, and ISO.

Workflow Automation

Trigger reviews, approvals, and reporting based on real-time events, with AI agents to help your team scale.

Official partner of the European Union

NeuralTrust has been recognized as the AI Security champion in the European Union, receiving the formal endorsement from the EU Commission

European Innovation Council

NeuralTrust

FAQ

How does NeuralTrust ensure compliance with AI regulations?

Our solutions help businesses comply with GDPR, HIPAA, PCI-DSS, and emerging AI regulations by enforcing data security policies, tracking AI interactions, and providing audit-ready reports to demonstrate regulatory compliance.

How does NeuralTrust improve AI security without impacting performance?

TrustGate is designed for low-latency processing, ensuring security without slowing down AI interactions. It analyzes traffic in real time, detects security risks, and enforces policies with minimal impact on response times.

How can I get started with NeuralTrust?

You can request a demo or speak with our team to explore how NeuralTrust can secure and optimize your AI operations. Contact us today to learn more.

Why choose NeuralTrust over any other alternative?

NeuralTrust offers the fastest AI gateway on the market (proven by benchmarks) ensuring secure, scalable, and low-latency AI operations unmatched by any alternative.

Does NeuralTrust support multi-cloud and on-premise AI deployments?

Yes. NeuralTrust is designed for multi-cloud environments and hybrid architectures. Whether you deploy AI on AWS, Azure, Google Cloud, or on-premises, our platform provides consistent security and observability across all environments.

Can each product be implemented separately?

Yes, each component of the NeuralTrust platform can be deployed independently, depending on your specific needs. However, for comprehensive protection and end-to-end visibility, we recommend implementing all three together to maximize security, control, and performance.

Start Implementing Secure AI Today With NeuralTrust

Ready to fortify your AI systems? Leverage NeuralTrust's unified platform to implement these mitigation strategies efficiently.

Sign Up

Book a Demo

LLM security standards we uphold

EU Artificial Intelligence Act

NIST

OWASP

MITRE ATLAS